

INTELLECTUAL PROPERTY LAW  
KIMBE, MARTENS, OLSON & B

A LIMITED LIABILITY PARTNERSHIP INCLUDING  
PROFESSIONAL CORPORATIONS

PATENT, TRADEMARK AND COPYRIGHT CAUSES

620 NEWPORT CENTER DRIVE

SIXTEENTH FLOOR

NEWPORT BEACH, CALIFORNIA 92660-8016

(949) 760-0404

FAX (949) 760-9502

INTERNET: WWW.KMOB.COM

LOUIS J. KNOBBE\*  
DON W. MARTENS\*  
GORDON H. OLSON\*  
JAMES B. BEAR  
DARRELL L. OLSON\*  
WILLIAM B. BUNKER  
WILLIAM H. NIEMAN  
ARTHUR S. ROSE  
JAMES F. LESNIAK  
NED A. ISRAELSEN  
DREW S. HAMILTON  
JERRY T. SEWELL  
JOHN B. SGANGA, JR.  
EDWARD A. SCHLATTER  
GERARD VON HOFFMANN  
JOSEPH R. RE  
CATHERINE J. HOLLAND  
JOHN M. CARSON  
KAREN VOGEL WEIL  
ANDREW H. SIMPSON  
JEFFREY L. VAN HOOSEAR  
DANIEL E. ALTMAN  
MARGUERITE L. GUNN  
STEPHEN C. JENSEN  
VITO A. CANUSO III  
WILLIAM H. SHREVE  
LYNDA J. ZADRA-SYMES\*  
STEVEN J. NATAUPSKY  
PAUL A. STEWART  
JOSEPH F. JENNINGS  
CRAIG S. SUMMERS  
ANNEMARIE KAISER  
BRENTON R. BABCOCK

THOMAS F. SMEGAL, JR.  
MICHAEL H. TRENHOLM  
DIANE M. REED  
JONATHAN A. BARNEY  
RONALD J. SCHOENBAUM  
JOHN R. KING  
FREDERICK S. BERRETTA  
NANCY WAYS VENSKO  
JOHN P. GIEZENTANNER  
ADEEL S. AKHTAR  
GINGER R. DREGER  
THOMAS R. ARNO  
DAVID N. WEISS  
DANIEL HART, PH.D.  
DOUGLAS G. MUEHLHAUSER  
LORI LEE YAMATO  
MICHAEL K. FRIEDLAND  
STEPHEN M. LOBBIN  
STACEY R. HALPERN  
DALE C. HUNT, PH.D.  
LEE W. HENDERSON, PH.D.  
DEBORAH S. SHEPHERD  
RICHARD E. CAMPBELL  
MARK M. ABUMERI  
JON W. GURKA  
ERIC M. NELSON  
MARK R. BENEDICT, PH.D.  
PAUL N. CONOVER  
ROBERT J. ROBY  
SABING H. LEE  
KAROLINE A. DELANEY  
JOHN W. HOLCOMB  
JAMES J. MULLEN, III, PH.D.

JOSEPH S. CIANFRANI  
JOSEPH M. REISMAN, PH.D.  
WILLIAM R. ZIMMERMAN  
GLEN L. NUTTALL  
ERIC S. FURMAN, PH.D.  
TIRZAH ABE LOWE  
GEOFFREY Y. IIDA  
ALEXANDER S. FRANCO  
SANJIVPAL S. GILL  
SUSAN M. MOSS  
JAMES W. HILL, M.D.  
ROSE M. THIESSEN, PH.D.  
MICHAEL L. FULLER  
MICHAEL A. GUILIANA  
MARK J. KERTZ  
RABINDER N. NARULA  
BRUCE S. ITCHKAWITZ, PH.D.  
PETER M. MIDGLEY  
THOMAS S. MCCLLENAHAN  
MICHAEL S. OKAMOTO  
JOHN M. GROVER  
MALLARY K. DE MERLIER  
IRFAN A. LATEEF  
AMY C. CHRISTENSEN  
SHARON S. NG  
MARK J. GALLAGHER, PH.D.  
DAVID G. JANKOWSKI, PH.D.  
BRIAN C. HORNE  
PAYSON J. LEMELLEUR  
WILLIAM G. BERRY  
DIANA W. PRINCE

OF COUNSEL

JERRY R. SEILER  
PAUL C. STEINHARDT

JAPANESE PATENT ATTY  
KATSUHIRO ARAI\*\*

EUROPEAN PATENT ATTY  
MARTIN HELLEBRANDT

KOREAN PATENT ATTY  
MINCHEOL KIM

SCIENTISTS & ENGINEERS  
(NON-LAWYERS)

RAIMOND J. SALENIEKS\*\*  
DANIEL E. JOHNSON, PH.D.\*\*  
JEFFERY KOEPKE, PH.D.\*\*  
KHURRAM RAHMAN, PH.D.  
JENNIFER A. HAYNES, PH.D.  
BRENDAN P. O NEILL, PH.D.  
THOMAS Y. NAGATA  
LINDA H. LIU  
YASHWANT VAISHNAV, PH.D.  
MEGUMI TANAKA  
CHE S. CHERESKIN, PH.D.\*\*  
ERIK W. ARCHBOLD  
PHILIP C. HARTSTEIN  
JULIE A. HOPPER  
CHRIS S. CASTLE  
JAMES W. AUSLEY  
R. P. CARON, PH.D.  
JENNIFER HAYES  
KIRK E. PASTORIAN, PH.D.  
CHARLES T. RIDGELY  
KEITH R. MCCOLLUM  
LANG J. MCHARDY

\* A PROFESSIONAL CORPORATION  
\* ALSO BARRISTER AT LAW (U.K.)  
\*\* U.S. PATENT AGENT

Assistant Commissioner for Patents  
Washington, D.C. 20231

**CERTIFICATE OF MAILING BY "EXPRESS MAIL"**

Attorney Docket No. : QLOGICP.021A

Applicants : Krantz, et al.

For : DISK CONTROLLER CONFIGURED TO  
PERFORM OUT OF ORDER EXECUTION  
OF WRITE OPERATIONS

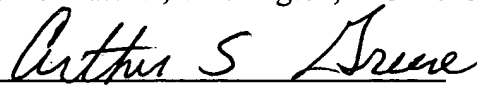
Attorney : Alexander Franco

"Express Mail"

Mailing Label No. : EL613492971US

Date of Deposit : August 22, 2000

I hereby certify that the accompanying Transmittal in Duplicate; Specification in 28 pages; 13 sheets of drawings; Declaration by Inventor in 2 pages; Recordation Form Cover Sheet and Assignment in 2 pages; Power of Attorney by Assignee in 2 pages; Checks for Filing Fees; and a Return Prepaid Postcard are being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and are addressed to the Assistant Commissioner for Patents, Washington, D.C. 20231.

  
ARTHUR GREENE

H:\DOCS\ASF\ASF-1559.DOC\dns  
082100

**DISK CONTROLLER CONFIGURED TO PERFORM OUT OF ORDER  
EXECUTION OF WRITE OPERATIONS**

5

Background of the Invention

Field of the Invention

10

This invention relates generally to hard disk controllers and, more particularly, the invention relates to a disk controller and associated methods that enable multiple write operations to be executed out of order during a single revolution of a disk.

Description of the Related Art

15

The speed at which a hard disk drive executes read and write operations requested by a host computer is critical to the performance of the computer. Any delay caused by the hard disk drive will likely cause a corresponding delay in the execution of a program by the computer. Some of the developments in hard drive technology have focused on decreasing the response time or increasing the data transfer rate for individual disk operations. Other developments in hard drive technology have focused on decreasing the amount of time that it takes to execute multiple operations. Two of the latter type of developments are described in U.S. Patents, 4,667,286 and 6,029,226.

20

25

U.S. Patent 4,667,286, titled "Method and apparatus for transferring data between a disk and a central processing unit," presents an architecture with toggling data buffers. The architecture allows multiple operations to be performed during a single revolution of the disk. Data is transferred between the central processing unit and a first data buffer while data is transferred between a second data buffer and the disk. The operation of the buffers can be toggled such that data is transferred between the central processing unit and the second data buffer while data is transferred between the first data buffer and the disk.

30

U.S. Patent 6,029,226, titled "Method and apparatus having automated write data transfer with optional skip by processing two write commands as a single write command," presents a method for combining write commands. The proximity of the

first logical block address (LBA) of a second write command is compared to the ending LBA of a first write command. If the second write command is sufficiently close, the disk controller executes the two write commands as a single write command, skipping over the sectors between the two commands during the combined write. Two writes, therefore, can be performed in one revolution of the hard disk.

In certain instances, it would also be advantageous to execute write operations in an order other than the order in which the operations are received by the disk controller. Executing write operations out of order may allow the writes to be executed during a single rather than multiple revolutions of the disk. The total time taken to execute the operations is therefore reduced. The present invention seeks to provide this advantage, among others.

#### Summary of the Invention

In the preferred embodiment, a hard disk unit includes a disk, a controller microprocessor, a host bus interface, a buffer memory, a buffer memory controller, and a disk formatter. Write operations are received by the bus interface, and the corresponding write operation data for each write operation is stored in a circular buffer in the buffer memory. The buffer memory controller implements the circular buffer and controls read and write access to the buffer memory. When write operation data arrives, the buffer memory controller stores the data contiguously within the circular buffer in the relative order in which the operation was received by the hard disk unit.

Upon receiving a command from the microprocessor, in the typical case, the buffer memory controller transfers the data of a write operation from the buffer memory to the disk formatter, which formats the data and writes it to the disk. In one embodiment, the buffer controller will continue transferring data from the circular buffer even after the data of the first write operation has been written to disk. The buffer controller supplies the data, operation by operation, in the order in which it was received by the hard disk unit. The microprocessor can simply issue commands to the disk formatter to continue writing the data of the subsequently received operations.

In the preferred embodiment, the buffer memory controller also includes a set of address registers and a set of block count registers. These registers allow the buffer

controller to transfer the data of several operations from the buffer memory in an order other than that in which the commands were received, using only a single command from the microprocessor.

In a preferred method, the microprocessor loads the address registers with the addresses within the buffer memory of the data of two or more write operations. The microprocessor loads the block count registers with the size, in blocks, of the corresponding operations. The microprocessor then issues a single command to the buffer memory controller to transfer the data identified by the address registers and block count registers from the buffer memory to the disk formatter. The buffer controller first transfers the data identified by the first address and block count register, then the second address and block count register, and so on. The address registers and block count registers enable the buffer memory controller to access the data of the write operations from the circular buffer in any order, regardless of the order in which the data is placed in the buffer.

The microprocessor is typically not fast enough to issue a second or additional commands to the buffer controller in time to supply the data for a second or additional operations without losing a revolution of the disk in the meantime. Typically, by the time the microprocessor sets up the buffer memory controller and issues the additional command, the sector to be written has passed the write head and a revolution of the disk is lost. In accordance with the preferred embodiment, however, two or more write operations to the same track may be completed during a single revolution of the disk, even if the operations are received in an order other than the order in which they must be written to the disk.

#### Brief Description of the Drawings

The present invention will be described below in connection with the attached drawings in which:

Figure 1 illustrates a typical configuration of a computer system in which a hard disk unit is utilized as a mass storage device of a host system;

Figure 2 illustrates a typical configuration of a hard disk controller;

Figure 3 illustrates the flow of data from the host system to the disk during a typical disk write operation;

Figure 4 illustrates the operation of a circular data buffer during write operations;

5 Figure 5 illustrates a buffer memory controller and a disk formatter in detail;

Figure 6 illustrates a method by which the hard disk controller typically handles write operations;

Figure 7 illustrates a schematic of a hard disk platter in conjunction with a write head;

10 Figure 8 illustrates a preferred embodiment of a buffer memory controller augmented to perform non-sequential reads of a buffer memory through a single command from a microprocessor

Figures 9A and 9B illustrate preferred methods performed by the buffer memory controller to achieve combined write functionality for write operations received out of order;

Figures 10A and 10B illustrate a preferred method performed by the hard disk controller to complete multiple out of order operations using a single set of commands to the buffer memory controller and the disk formatter; and

20 Figure 11 illustrates a method of an example sequence of steps performed by the hard disk controller for two write operations that arrive out of order.

#### Detailed Description of the Embodiments

In the following description, reference is made to the accompanying drawings, which form a part hereof, and which show, by way of illustration, specific embodiments or processes in which the invention may be practiced. Where possible, the same reference numbers are used throughout the drawings to refer to the same or like components. In some instances, numerous specific details are set forth in order to provide a thorough understanding of the present invention. The present invention, however, may be practiced without the specific details or with certain alternative equivalent components and methods to those described herein. In other instances, well-

known methods and components have not been described in detail so as not to unnecessarily obscure aspects of the present invention.

## I. HARD DISK CONTROLLER OPERATION

### A. Computer System

Figure 1 illustrates a typical configuration of a computer system in which a hard disk unit 110 is utilized as a mass storage device of a host system 120. The host system 120 is preferably the motherboard of the computer system and typically includes a system bus 122 that connects one or more microprocessors, RAM, and possibly other components. The hard disk unit 110 typically includes a hard disk controller 112 that controls a disk 114. The disk 114 typically includes several platters of rotating magnetic media that are read from and written to by read and write heads (not illustrated). In the illustrated embodiment, the hard disk controller 112 is connected to the host system through an Advanced Technology Attachment (ATA) bus 116. The ATA bus 116 is in turn connected to the host system bus 112 through an I/O bus bridge adapter 124. The adapter 124 serves as a bridge between the different protocols and technologies of the host system bus 122 and the ATA bus 116. Alternative technologies, such as SCSI, USB, or Firewire may also be used in place of the ATA bus 116.

The host system 120 typically sends read and write operations to the hard disk controller 112 via the ATA bus 116. A read operation typically identifies the requested data by the location on the disk of the beginning logical block address (LBA) of the data. The LBA specifies the head, track, and sector number of a block (512K) of data. A sector typically holds one block of data. The request causes the hard disk unit 110 to respond by reading and transmitting the requested data back to the host system 120 via the ATA bus 116. A write operation typically identifies the starting LBA of the location on the disk to be written. Data is transferred from the host system 120 to the hard disk controller 112 and is written to the disk starting at the specified LBA.

### B. Hard Disk Controller

Figure 2 illustrates a typical configuration of a hard disk controller 112. The hard disk controller 112 typically includes a microprocessor 204, a set of firmware code 205, a buffer memory 206, and a hard disk controller core 210. The controller core 210

handles the flow of data between the ATA bus 116 and the disk 114. The controller's tasks typically include transferring data between the disk 114 and the buffer memory 206, transferring data between the buffer memory 206 and the ATA bus 116, and performing error correction and CRC calculations. The operation of the controller core 210 is in turn controlled by the controller microprocessor 204, which typically executes the set of firmware code 205. The firmware code 205 may be stored in nonvolatile memory. The microprocessor may also have an additional interface (not illustrated) to the disk 114 to handle tasks such as moving the read/write heads of the disk 114 to the proper track. The buffer memory 206 is used to buffer data between the disk 114 and the ATA bus 116 to compensate for delays, latency, and timing differences between the two devices. In some configurations, the buffer memory 206 also serves as the memory for the controller microprocessor 204. The size of the buffer memory 206 may range from 256 kilobytes to 8 megabytes, depending upon the application.

The controller core 210 typically includes an ATA bus interface 220, a buffer memory controller 230, and a disk formatter 240. The ATA bus interface 220 implements the ATA bus protocol to receive write operations and read operations and to send read data to the host system 120. A write operation 222 is shown in dashed lines to indicate that it passes through the bus interface 220. The LBA of the write operation is typically read by the microprocessor 204 and stored. The bus interface 220 passes the data to the buffer memory 206 through the buffer memory controller 240.

The buffer memory controller 230 controls interleaved access to the buffer memory 206 by the ATA bus interface 220, the disk formatter 240, and the microprocessor 204. The control of the buffer memory 206 by the buffer memory controller 230 typically involves arbitrating access by the various units as well as buffering data, yet again, between the buffer memory 206 and the units that access it. The buffer memory controller 230 typically includes an ATA bus interface data FIFO (first in, first out buffer) 232 that buffers data between the ATA bus interface 220 and the buffer memory 206, a disk formatter data FIFO 234 that buffers data between the buffer memory 206 and the disk formatter 240, and buffer controller logic 236 that controls the functionality of the buffer controller 230. The buffer controller 230 typically stores write operation data in a circular data buffer 208 that the buffer

controller 230 implements within the buffer memory 206, preferably using pointers. The ATA bus interface data FIFO 232 and the disk formatter data FIFO 234 are relatively small FIFO buffers, preferably about 64 bytes each. The circular data buffer 208, on the other hand, is preferably much larger, possibly occupying the whole buffer memory 206, which may be 256 kilobytes to 8 megabytes in some embodiments.

The disk formatter 240 controls the actual writing of data to the disk 114. The disk formatter 240 typically receives data from the buffer memory 206 through the buffer memory controller, formats the data for writing to the disk, possibly adding error correction codes, and sends the data to the write head of the disk. The disk formatter 240 also monitors the sector of a track over which the disk's read/write head is positioned to determine the proper timing for sending the data to the write head. The functionality of the disk formatter is controlled by disk formatter logic 242.

Figure 3 illustrates the flow of data from the host system 120 to the disk 114 during a typical disk write operation. The data is first transferred from the host system 120 via the ATA bus 116 to the ATA bus interface 220 of the hard disk controller 112. The ATA bus interface 220, in turn, transfers the data through the ATA bus interface FIFO 232 of the buffer memory controller 230 to the circular data buffer 208 in the buffer memory 206. The data remains in the circular data buffer 208 until it is to be written to the disk 114. If the disk heads have been correctly positioned and the microprocessor 204 has determined that the data is ready to be written to disk 114, the microprocessor 204 issues a command to the buffer controller 230 to begin transferring the data into the disk formatter data FIFO 234. The microprocessor 204 also issues a command to the disk formatter 240 to begin reading the data from the disk formatter data FIFO 234, formatting the data, and sending the data to the write head of the disk 114. A disk formatter command typically only instructs the disk formatter 240 to write data to one track on the disk. In the case of a large write operation, for example, on the order of several megabytes, earlier data already placed in the circular data buffer 208 is read by the disk formatter 240 while later data for the same operation is being written to the circular data buffer 208 by the ATA bus interface 220. If the write operation involves writing to multiple tracks, after the disk formatter 240 writes data to each track,



the microprocessor 204 moves the write head to the next track and issues an additional command to the disk formatter 240 to write data to the next track.

Figure 4 illustrates the operation of the circular data buffer 208 during write operations. In the illustrated example, the data of three write operations have been written to the buffer memory 206 by the ATA bus interface 220. The buffer memory controller 230 updates a write pointer 402 to indicate the address within the buffer memory 206 at which the controller writes incoming data from the ATA bus interface 220. As data is written, the pointer 402 is updated. If the write pointer 402 reaches the end of the circular buffer 208, the pointer 402 is wrapped around to the other end to effect the circular buffer functionality. A read pointer 404 indicates the address at which data is read from the circular buffer 208 to be transferred to the disk formatter 240. In the illustrated example, the data of the three write operations in the buffer 208 has not yet been transferred to the disk formatter 240. As data is read, the default operation of the buffer memory controller logic 236 advances the read pointer 404 within the circular buffer 208. The read pointer is eventually wrapped around like the write pointer. In one embodiment, as long as the buffer controller is not reset, the read pointer 404 will continually advance through the data of write operation after write operation as long as the disk formatter 234 continues to read the data from the disk formatter data FIFO 234. The write pointer 404 typically will not, however, advance past the read pointer 402.

Figure 5 illustrates the buffer memory controller 230 and the disk formatter 240 in greater detail. The buffer memory controller 230 includes a disk formatter (DF) buffer address register 502 in which the read pointer 404 is maintained during write operations. Typically, the buffer address register 502 is initially loaded by the microprocessor 204. The buffer memory controller logic 236, however, automatically advances the buffer address register 502 as data is transferred to the disk formatter data FIFO 234.

The disk formatter 240 typically includes a sector requested register 512, a sector count register 514, and a set of skip sector registers 516. The microprocessor 204 initiates a disk operation by loading the sector requested register 512, the sector count register 514, and the skip sector registers 516 with the appropriate values and by issuing



memory. The microprocessor 204 may, in this case, store the LBA in the buffer memory 206, accessing the buffer memory through the buffer memory controller 230. Before control proceeds to the remaining steps of the method 600, additional write operations may be received by the ATA bus interface 220 and handled by the microprocessor 204 as indicated by the arrow leading from the step 606 to the step 602.

At a step 608, the microprocessor 204 decides to execute a received write operation. The microprocessor 204 is preferably programmed by the firmware code 205 in accordance with known techniques in order to determine when to execute the operation.

At a step 610, the microprocessor 204 loads the disk formatter buffer address register 502 with the beginning address of the write operation data within the buffer memory 206. At a step 612, the microprocessor 204 completes programming of the buffer controller 204 for the write operation. This step may involve clearing the disk formatter data FIFO 234 of any stale data and/or loading a command into a command register of the buffer memory controller 230.

At a step 614, the buffer memory controller 230 begins to transfer data from the buffer memory 206 to the disk formatter data FIFO 234. As the disk formatter 240 reads the data from the FIFO 234, the buffer memory controller 230 writes data to the FIFO 234 to keep the FIFO 234 from underflowing.

At a step 616, the microprocessor 204 programs any additional components necessary to move the write head of the disk 114 to the proper track. The circuits that control the movement of the write head are preferably external to the controller core 210 but may be included in the core 210. The step 616 may alternatively be performed between the steps 610 and 608.

At a step 618, the microprocessor 204 loads the sector requested register 512 of the disk formatter 240 with the sector ID of the first sector of the write operation to be written to the track. At a step 620, the microprocessor 204 loads the sector count register 514 with the number of sectors to be written to the track. In the case a write operation spans multiple tracks, the number of sectors that are written to the track may be less than the number of blocks in the entire operation. In this case, the disk formatter

240 is later programmed to write data to the subsequent tracks after the write to the present track has been completed.

At a step 622, the microprocessor 204 loads the skip sector registers 516 with data identifying defective sectors on the track to be skipped during the write to the track.

5 At a step 624, the microprocessor 204 completes programming of the disk formatter 240 for the write to the track. This step typically includes loading a command into a command register of the disk formatter 240.

10 At a step 626, the disk formatter 240 begins execution of the command loaded into its command register by the microprocessor 204. The disk formatter 240 reads data from the disk formatter data FIFO 234, formats the data for writing to the disk 114, and sends the data to the write head of the disk 114. The formatting of the data for writing to the disk 114 may involve adding a cyclic redundancy check (CRC) code to the data. The disk formatter 240 typically also determines the proper timing for sending the data to the write head of the disk 114 so that the data is written to the proper location on the track.

15 At a step 628, the disk formatter 240 completes the writing of the last sector to the current track. At a step 630, the disk formatter 240 sends an interrupt to the controller microprocessor 204 and indicates the success status of the write to the track. Typically the write will have been successful. If the write was unsuccessful, the microprocessor with then handle the problem.

20 At a step 632, the microprocessor determines whether there is more data to be written to another track in conjunction with the present write operation. If additional data remains to be written to another track, control passes back to the step 616. From the step 616, the microprocessor moves the write head to the next track and programs the disk formatter 240 to write more data to the next track.

25 At the step 632, if the current write operation has completed because no more data remains to be written to another track, the hard disk controller 112 is ready to handle the receipt or execution of another write operation. Accordingly, control passes to either the step 602 or the step 608, possibly depending upon whether another write operation has already been received.

30

#### D. Combining Write Operations to Improve Performance

In one embodiment, the microprocessor 204 may be programmed to complete two or more write operations with only a single command to the buffer memory controller 230. The buffer memory controller need not be reprogrammed as in the steps 610 – 614 as long as the data from the second write operation is written contiguously after the data of the first write operation within the circular data buffer 208 before the first operation completes. The buffer memory controller 230 is preferably configured to continue supplying data from circular data buffer 208 to the disk formatter data FIFO 242 as long as the disk formatter 240 keeps reading the data. Accordingly, the buffer memory controller 230 will continue to supply data to the disk formatter data FIFO 234 from the circular data buffer 208 as the last of the data of the first write operation is read by the disk formatter 240 and written to disk. The data that the buffer memory controller 230 automatically supplies to refill the FIFO 234 will be the data of the second write operation. In this instance, therefore, the buffer memory controller 230 need not be reprogrammed, and the steps 610 – 614 of the process 600 can be skipped.

In one embodiment, the microprocessor 204 may also be programmed to partially or completely write data from two write operations using a single command to the disk formatter 240. A single command to the disk formatter can be used when the ending LBA of an earlier operation has the same track but a lower sector number than the beginning LBA of the subsequent operation. In this case the microprocessor can treat the two operations as a single operation by instructing the disk formatter 240 to skip the sectors between the two operations. This can be accomplished by loading one or more of the skip sector registers 516 with the registers to be skipped between the write operations.

The aforementioned techniques allow the programming of the buffer controller 230 and/or the disk formatter 240 for two or more write operations to be combined. One advantage to combining the programming of the disk formatter 240 for two or more write operations is that the write operations will be completed in a single revolution.

The microprocessor 204 typically has a much longer response time and is much slower than the disk formatter 240 to react. There is typically not enough time after the disk formatter 240 interrupts the microprocessor 204 upon completion of a write

operation to allow the microprocessor 204 to reprogram the disk formatter 240 for a second write operation without losing a revolution of the disk 114. Accordingly, in order to complete more than one write operation during a single revolution of the disk 114, the two or more write operations must be completed through a single command to the disk formatter 240. In order to combine the programming of the disk formatter 240 for the two or more operations the programming of the buffer controller 230 for these operations must also be combined. This is the case since the microprocessor 204 is also typically not fast enough to reprogram the buffer controller 230 between write operations.

As already mentioned, the controller microprocessor 204 is generally not fast enough to reprogram the buffer memory controller 230 or the disk formatter 240 between the write operations. In addition, some disk formatters may be configured to begin execution of each command by searching for the first sector on a track, such as, for example, to begin a timing sequence for writing the data. In such cases, the disk formatter 240 will nevertheless be able to execute only one command per revolution, regardless of the speed of the microprocessor 204. Accordingly, if several write operations that write data to the same track on a disk are to be executed in a single revolution, the write operations must typically be completed through a single command to each of the buffer controller 230 and the disk formatter 240.

Executing more than one write operation in a single revolution of the disk 114 substantially speeds up the execution of the writes. For example, if two writes are executed in a single revolution, execution time can be decreased to as little as 1/2 of the time of separately executing the writes. If four writes are executed in a single revolution, execution time will be at most 1/3 and as little as 1/4 of the time of executing the writes separately.

#### E. Limitations of the Aforementioned Techniques

Although the aforementioned techniques can be used to decrease the time to perform several write operations in certain instances, in other instances, these techniques cannot be used. For example, if the order in which the write operations arrive does not correspond to the order of the respective write operations on the disk, these techniques cannot be used.



A)

block count registers 862, a working block count register 864, a busy flag 854, and combined write logic 870.

There are preferably 8 or more registers in the set of address registers 852. Although 4, 3, or 2 address registers could be used, this would limit the number of write operations that could be written through a single set of commands to the buffer memory controller 830 and the disk formatter 240. The microprocessor 204 loads each of the address registers 852 with the starting address of the data in the buffer memory 206 for one of the write operations to be written to the track. The address registers 852 are loaded in the order in which the write operations are to be performed. The address registers 852 are preferably configured to be written to by the microprocessor 204 through a FIFO-like interface, similar to the skip sector registers 516. Each of the addresses in the set of address registers 852 is eventually transferred into the working disk formatter buffer address register 802 during execution of the write operations. The working address register 802 functions substantially in the same way as the buffer address register 502.

The number of block count registers 862 matches the number of address registers 852. The microprocessor 204 loads each of the block count registers 862 with the number of blocks of data stored in the buffer memory 206 to be written to the track for the corresponding write operation. The block count registers 862 are loaded in the order in which the write operations are to be performed. Like the address registers 852, the block count registers 862 are preferably configured to be written to by the microprocessor 204 through a FIFO-like interface, similar to the skip sector registers 516. Each of the block counts in the block count registers 862 is eventually transferred into the working block count register 864 during the execution of the write operations. The working block count register 864 is used to count down the number of blocks remaining to be transferred as they are transferred from the buffer memory 206 to the disk formatter data FIFO 234 for each operation.

The busy flag 854 is preferably a register or a bit in a register that can be read by the microprocessor 204 to determine whether the address registers 852 and the block count registers 862 are full. If the register sets 852 and 862 are full, the bit is set. If the register sets 852 and 862 are not full, the bit is cleared.



The combined write logic 870 is preferably added to the buffer memory controller logic 236 to control the functionality of the block count registers 862, the address registers 852, the working block count register 864, the busy flag 854. The logic 870 controls these elements to transfer the data of multiple write operations to the disk formatter data FIFO 234 from a single command supplied by the microprocessor 204 regardless of the order of the operations' data in the buffer memory 206.

The disk formatter 240 of the preferred embodiment is substantially unmodified. The number of skip sector registers 516, however, are preferably increased. In the preferred embodiment, 16 skip sector registers 516 are included instead of 8 to handle the extra sector skips between the two or more write operations on the same track.

#### B. Buffer Memory Controller Methods

Figures 9A and 9B illustrate preferred methods 900 and 950 performed by the controller 830 under the control of the combined write logic 870 and the buffer memory controller logic 236 to achieve the combined write functionality for write operations received out of order. The methods 900 and 950 may be performed simultaneously, preferably in an interleaved fashion. The method 900 sets the busy flag 854, while the method 950 manipulates the working buffer address 802 and working block count 864 registers.

At a step 902 of the method 900, the buffer memory controller 830 receives data in the address registers 852 and the block count registers 862. At a step 904, the controller 830 determines whether the address and block count registers 852 and 862 are all full. If the registers 852 and 862 are full, the controller 830 sets the busy flag 854. If the registers 852 and 862 are not full, the controller 830 passes control back to the step 902.

At a step 952 of the method 950, the buffer memory controller 830 receives a command from the microprocessor 204 that selects the use of the address registers 852 and the block count registers 862 to transfer the data of multiple operations. The use of these registers 852 and 862 is preferably made an option to the microprocessor 204, the option may only be selected when the operations to be combined have arrived out of order.

At a step 954, the buffer memory controller 830 transfers an address of data for a write operation from the address registers 852 into the working address register 802. The addresses are transferred from the address registers 852 in the order in which the address registers 852 were loaded by the microprocessor 204. In other words, the address registers 852 preferably have a FIFO functionality. At a step 956, the controller 830 transfers a block count for a write operation from the block count registers 862 into the working block count register 864. The block count registers 862 also have a FIFO functionality. At a step 858, the controller clears the busy flag since at least one location in each of the sets of registers 852 and 862 has just been emptied.

At a step 960, the controller 830 begins transferring data from the buffer memory 206 to the disk formatter data FIFO 234, starting at the address in the working address register 802. The controller preferably uses the working address register 802 as a read pointer and increments the register 802 as it reads data from the buffer memory 206.

At a step 962, the controller 830 decrements the working block count register 864 upon initiating the transfer of each block (512K) of data from the buffer memory 206 to the disk formatter data FIFO 234.

At a step 964, the controller 830 waits for the working block count register 864 to reach zero. This indicates that the last block of the operation has begun to be transferred from the buffer memory 206 to the disk formatter data FIFO 234. At a step 966, the controller 830 transfers the last block of data for the operation and then stops transferring data. At this point control passes back to the step 954, and the data of the next write operation is transferred.

The process 950 continues until all of the entries in the address registers 852 and the block count registers 862 have been processed. Once all of the entries have been processed, the controller 830 is ready to receive another command from the microprocessor 204.

#### C. Hard Disk Controller Method

Figures 10A and 10B illustrate a preferred method 1000 performed by the hard disk controller 112 to complete multiple out of order operations using a single set of commands to the buffer memory controller 830 and the disk formatter 240. The

microprocessor 204 is preferably programmed with the firmware 205 to perform the method 1000. The steps 1002, 1004, and 1006 are substantially the same as the steps 602, 604, and 606 of the method 600.

5       At a step 1008, the microprocessor 204 identifies two or more write operations that write data to the same track, where one or more of the operations arrive out of order with respect to the locations of the corresponding sectors on the track. At a step 1010, the microprocessor 204 orders the operations based upon the locations of their respective sectors on the track.

10       At a step 1012, the microprocessor 204 loads the address registers 852 with the beginning addresses in the buffer memory 206 of the data blocks of the write operations. In the preferred embodiment, the address registers 852 are written through a FIFO-like interface, which maintains the order in which the addresses are written. Accordingly, the addresses of the data blocks are loaded in the order, (determined by the microprocessor 204 in the step 1010) in which the corresponding write operations are to  
15       be written to the track. At a step 1014, the microprocessor loads the block count registers 862 with the corresponding number of data blocks for each write operation in the same manner.

20       At a step 1016, the microprocessor 204 completes programming the buffer memory controller 830. The programming preferably includes issuing a command to the buffer memory controller 830 instructing it to transfer data based upon the addresses and block counts loaded in the address registers 852 and the block count registers 862. The microprocessor 204 preferably also has the option of programming the buffer memory controller 830 to perform a regular data transfer in accordance with the method 600.

25       At a step 1018, the buffer memory controller 830 begins to transfer data of the write operations from the buffer memory 206 to the disk formatter FIFO 234. As the disk formatter 240 reads the data from the FIFO 234, the buffer memory controller 830 writes data to the FIFO 234 to keep the FIFO 234 from underflowing.

30       At a step 1020, the microprocessor 204 programs any additional components necessary to move the write head of the disk 114 to the proper track. The step 1020 may alternatively be performed between the steps 1012 and 1010.

At a step 1022, the microprocessor 204 loads the sector requested register 512 of the disk formatter 240 with the sector ID of the first sector of the first write operation to be written to the track. At a step 1024, the microprocessor 204 loads the sector count register 514 with the total number of sectors to be written to the track for all of the identified operations.

At a step 1026, the microprocessor 204 loads the skip sector registers 516 with data identifying sectors on the track to be skipped between write operations. The registers are also loaded with the defective sectors on the track to be skipped during the writes. At a step 1028, the microprocessor 204 completes programming of the disk formatter 240 for the write to the track. This step typically includes loading a command into a command register of the disk formatter 240.

At a step 1030, the disk formatter 240 begins execution of the command loaded into its command register by the microprocessor 204. The disk formatter 240 reads data from the disk formatter data FIFO 234, formats the data for writing to the disk 114, and sends the data to the write head of the disk 114.

At a step 1032, the buffer memory controller 830 supplies data from the successive write operations according to the addresses and block counts loaded in the address registers 852 and the block count registers 862. The buffer memory controller 830 preferably performs the method 950 in this step.

At a step 1034, the disk formatter 240 completes the writing of the last sector to the current track. At a step 1034, the disk formatter 240 sends an interrupt to the controller microprocessor 204 and indicates the success status of the write to the track.

In the case that all of the write operations are completely on a single track, the method ends at this point and the hard disk controller 112 is prepared to perform additional write operations. Alternatively, the method 1000 may also be used when the first write operation and/or the last write operation to be executed span two or more tracks. For example, a short write operation to a track may arrive before a longer operation that ends on the same track. The latter operation can be executed first and followed by the first operation using a single buffer memory controller command. A single disk formatter command can be used to write the last portion of the latter

command's data and all of the former command's data. In the case one or more write operations span two tracks, the method 1000 continues as described below.

At a next step 1036, the microprocessor determines whether there is more data to be written to another track in conjunction with the present set of write operations. If additional data remains to be written to another track, control passes back to the step 1020. From the step 1020, the microprocessor moves the write head to the next track and programs the disk formatter 240 to write more data to the next track.

At the step 1038, if the current set of write operations has completed since there remains no more data to be written to another track, the hard disk controller 112 is ready to handle the receipt or execution of another set of write operations. Accordingly, control passes to either the step 1002 or the step 1008, possibly depending upon whether another write operation has already been received.

#### D. Example Operation

Figure 11 illustrates a method 1100 of an example sequence of steps performed by the hard disk controller 112 for two write operations that arrive out of order. The method 1100 illustrates in more specific detail for the case of two write operations, the steps 1002 through 1014 of the method 1000.

At a step 1102, the ATA bus interface 220 receives a first write operation. At a step 1104, the buffer memory controller 230 writes the data of the first write operation to the circular data buffer 208. At a step 1106, the bus interface 220 receives a second write operation. At a step 1108, the buffer memory controller 230 writes the data of the second write operation to the circular data buffer 208 after the data of the first write operation. The data of the second write operation may but need not be contiguous with the data of the first write operation in the buffer memory 206.

At a step 1110, the microprocessor 204 determines that the second write operation writes to the same track as the first write operation. The microprocessor 204 preferably also determines that there are no commonly written sectors between the two write operations to avoid creating stale data. If the two operations involve a common sector, then the operations should be written in the order in which they arrive.

At a step 1112, the microprocessor 204 determines that the second write operation is located before the first write operation on the track. The microprocessor

preferably makes this determination by comparing the sector numbers of the first sector that each operation writes. The order of the operations is determined relative to the position where the write head is capable of first writing to the track. In some embodiments, the disk formatter 240 may be configured to be able to first write data to any sector on the track. If the head arrives at a track with 1024 sectors and stabilizes itself while the head is over the 511th sector, for example, these embodiments may be configured to be able to write data to the 512th sector during the current revolution of the disk. In this case, if the first write operation writes to sector 127 and the second operation writes to sector 575, the second write operation is located before the first write operation on the track. In other embodiments, the disk formatter 240 may be configured to wait for and detect the 0th sector on a track before writing any data. In these embodiments, even if the head stabilizes itself over the 511th sector of a track with 1024 sectors, the disk formatter must wait until the 0th sector passes below the head before any data is written. In this case, if the first write operation writes to sector 575 and the second write operation writes to sector 127, the second write operation is located before the first write operation on the track.

In one embodiment, the steps 1110 and 1112 can be combined by comparing the starting and ending sector numbers of each write operation. In this case, both operations must also write to the same track. In addition, supposing the disk formatter 240 must wait for the 0th sector to pass before writing data, the ending sector number of the second operation must be less than the starting sector number of the first operation. It will be apparent to one skilled in the art how to handle the case where the disk formatter 240 need not wait for the 0th sector to pass before beginning an operation.

At a step 1114, the microprocessor 204 loads one of the address registers 852 with the address of the data of the second write operation in the circular data buffer 208. At a step 1116, the microprocessor 204 loads one of the block count registers 862 with the number of data blocks in the second write operation. At a step 1118, the microprocessor 204 loads one of the address registers 852 with the address of the data of the first write operation in the circular data buffer 208. At a step 1120, the microprocessor 204 loads one of the block count registers 862 with the number of data blocks in the first write operation.

Once the buffer memory controller 830 has been loaded with the addresses and block counts of the two operations, the method 1100 continues at the step 1016 of the method 1000.

#### E. Extensions and Additional Embodiments

The method 1100 can be generalized by one skilled in the art into the firmware code 205 to handle several write operations. The generalized firmware 205 can be configured to program the buffer memory controller 830 to provide data of multiple write operations to the disk formatter 240 regardless of the order in which the operations are received.

The number of operations that can be handled through a single command to the buffer memory controller 830 may be limited by the number of registers in the address registers 852 and the block count registers 862. In the preferred embodiment, however, these sets of registers 852 and 862 can be written by the microprocessor 204 while the buffer memory controller 830 is executing a command. This functionality allows a potentially infinite number of out of order write operations to be handled with a single command to the buffer controller 830.

#### III. CONCLUSION

Although the invention has been described in terms of certain preferred embodiments, other embodiments that are apparent to those of ordinary skill in the art, including embodiments which do not provide all of the features and advantages set forth herein, are also within the scope of this invention. Accordingly, the scope of the invention is defined by the claims that follow. In the claims, a portion shall include greater than none and up to the whole of a thing. In the method claims, reference characters are used for convenience of description only, and do not indicate a particular order for performing the method.